

Symmetry-factored Statistical Modelling of Craniofacial Shape

Hang Dai, William A. P. Smith and Nick Pears

Department of Computer Science, University of York, UK

{hd816,william.smith,nick.pears}@york.ac.uk

Christian Duncan

Alder Hey Hospital, Liverpool, UK

Christian.Duncan@alderhey.nhs.uk

Abstract

We present a new method for symmetry-factored statistical modelling of 3D shape. Our method comprises three novel components. First, a means to symmetrise a 3D mesh, regularised using the Laplace-Beltrami operator. Second, a symmetry-aware variant of Generalized Procrustes Analysis (GPA). Third, a means to compute a linear statistical shape model in which symmetry and asymmetric shape variation are modelled separately. We focus on human head data and build the first 3D morphable model of craniofacial asymmetry. The qualitative and quantitative evaluation demonstrates that the proposed model outperforms a linear model that does not decompose symmetric and asymmetric variation. It also validates that symmetry-aware GPA can improve the data generalisation and reconstruction ability of the standard PCA model. We will make our model and the implementation of our method publicly available¹.

1. Introduction

Most biological objects, including human heads and bodies, possess approximate symmetries. Often this is principally extrinsic, bilateral symmetry (i.e. reflective symmetry about a vertical plane bisecting the object). Deviations from exact symmetry are an interesting and potentially important geometric property in terms of modelling and understanding 3D shape variation. The degree of asymmetry may convey information about an object. For example, there is evidence that facial asymmetry is used by humans to measure genetic health [14] and asymmetry in man-made objects may indicate imperfections in the manufacturing process. For this reason, we suggest that statistical shape models should separate symmetric from asymmetric shape variability in order to reveal interesting and potentially subtle aspects of shape variation over population datasets.

The asymmetric variation in human head shape is much smaller than symmetric variation. Hence, classical models obtained by applying Principal Component Analysis

(PCA) directly to head shapes pools asymmetric variation with the much larger symmetrical variation. In this paper, we consider how to build a symmetry-factored statistical shape model in a principled way. Once built, the symmetry-factored statistical modelling of craniofacial shape opens up new opportunities not afforded by existing 3D morphable models (3DMMs). Firstly, it provides a different way to reconstruct a full head model from a profile scan beyond just using the reflection of the profile. Secondly, we can analyze the asymmetry of multiple craniofacial regions. Thirdly, we can investigate the different importance of symmetry-factored information in different applications, e.g. face recognition or ergonomics.

Recently we developed a fully automatic pipeline for building 3DMMs [9]. We did not exploit symmetry and asymmetry constraint, which is the primary concern of this paper. Here, our main contribution is to show how to build a statistical model with separate parameters for symmetric and asymmetric variations. The resulting model is still linear and so can be used in place of any existing 3DMM but with the additional ability to separate symmetric from asymmetric variation. Secondary contributions include a method for symmetrisation regularised by the Laplace-Beltrami operator, symmetry-aware Generalized Procrustes Analysis (GPA) and the symmetry-factored statistical modelling method. Tertiary contributions include making available, for the first time, a morphable model of craniofacial asymmetry. Comprehensive evaluation shows that the proposed model has significantly better performance than the linear models.

2. Related work

There is a large body of work on detecting symmetries in images [28], 2D shapes, 3D meshes and point clouds. These symmetries could be exact or approximate [22], extrinsic [30] or intrinsic [25], partial [28, 34], hierarchical or full [25]. Once detected, this enables symmetry-aware mesh processing [12] to take place in which symmetries help regularise and de-noise a range of processing tasks.

We are not the first to consider symmetry in the context of statistical shape modelling. The geometric morphomet-

¹<https://www-users.cs.york.ac.uk/~nep/research/LYHM/>

rics community [20, 15, 16, 17] have built models of deviations from symmetry, though this has largely been in 2D. Savriama et al. [33] present a decomposition of asymmetric shape into a symmetric shape and asymmetry variation. There are many works on how to measure facial asymmetry variation. There is much literature from both the Computer Vision perspective [19, 5, 21, 37] and the Biology perspective [6, 13, 16]. Symmetrisation is crucial to the acquisition of asymmetry variation [23]. The literature closest to our work is that of Liao et al. [18]. However, this work has three limitations: i) the authors applied a linear average of the left and right sides of the face, which is the basic version of symmetrisation, to force the shape to be symmetric; ii) their symmetrisation relies on facial landmarks, which limits the usage of the method. iii) there was no statistical modelling to study variation in facial symmetry and asymmetry. In this paper, we pose symmetrisation as finding the symmetric mesh that minimises distortion relative to a given mesh, where distortion is measured via the Laplace-Beltrami operator. Moreover, our method does not require landmarks. Finally, we build a symmetry-factored statistical model to analyze both symmetrical and asymmetrical variation.

Blanz and Vetter built a 3D morphable model (3DMM) from 3D face scans [2] and employed it in a 2D face recognition application [3]. Two hundred scans (100 males and 100 females) were used to build the model. Dense correspondences were computed using optical flow with an energy that depended on both shape and texture. Currently, the Basel Face Model (BFM) is the most well-known and widely used 3D face model and was developed by Paysan et al. [29]. Again 200 scans were used, but the method of determining corresponding points was improved. Instead of optical flow, a set of hand-labelled feature points is marked on each of the 200 training scans. The corresponding points on a template mesh are known, which is then morphed onto the training scan using under-constrained per-vertex affine transformations, which are constrained by regularisation across neighbouring points. The technique is known as optimal-step Non-Rigid Iterative Closest Points (NICP) [1]. Booth et al. [4] built a Large Scale Facial Model (LSFM), using the same NICP template morphing approach, followed by GPA for alignment, and PCA for the model construction. Recently Hang et al. [9] used a self-adapting dense correspondence establishing system combining a hierarchical parts-based template morphing framework in the shape channel and a refining optical flow in the texture channel to build the first 3DMM of the full human head including the cranium. These 3DMMs are either standard PCA linear models or multi-linear models, but there are no 3DMMs of craniofacial asymmetry.

The previous 3DMMs employ deformable template methods [29, 10, 4, 32, 31] or optical flow [2, 8] to establish dense correspondence. When applied to 3D craniofacial

data, these methods have either over-fitting or under-fitting problem, which makes it unfeasible to find the correct symmetric partner. This is validated in [9] and we apply the same template morphing method as that paper. This hierarchical parts-based template morphing framework is able to avoid both over-fitting and under-fitting and we exploit this in the pipeline proposed here.

Many works [36, 7, 26, 27, 24] addressed the problem of finding the facial symmetry plane or contour. But the symmetry plane or contour problem is not the focus of this paper. We apply a standard method to extract the symmetry plane from 3D scans.

3. Symmetry-factored modelling

In symmetry-aware geometry processing, the task of symmetry detection is usually decoupled from that of symmetry-aware processing. In our proposed symmetry-factored statistical modelling pipeline, we follow this idea with a further decomposition into five processes that we treat independently:

- (i) **Symmetry detection and symmetric remeshing:** Given an arbitrary mesh, for every vertex attempt to find a corresponding symmetric partner vertex or point on the surface. Then remesh the surface with symmetric topology so that every vertex has a symmetric partner or is self-symmetric.
- (ii) **Symmetry plane alignment:** Given a mesh with symmetric topology, estimate a plane which maximises the symmetry about the plane.
- (iii) **Symmetrisation:** Given a mesh with symmetric topology, find the “closest” shape with exact geometric symmetry.
- (iv) **Symmetry-aware GPA:** Given a collection of meshes with symmetric topology with symmetry planes aligned, solve for the remaining degrees of freedom to best align the collection.
- (v) **Symmetry-factored statistical modelling:** Separately model symmetric and asymmetric modes of shape variation.

Our specific aim in this paper is to model deviations from exact extrinsic symmetry for a set of 3D head scans with neutral facial expression.

4. Preliminaries

We represent a mesh either as a matrix of vertex positions $\mathbf{X} \in \mathbb{R}^{N \times 3}$ or as a vector $\mathbf{x} \in \mathbb{R}^{3N} = \text{vec}(\mathbf{X}) = [x_1, y_1, z_1, \dots, z_N]^T$. The i th vertex in the mesh, $\mathbf{v}_i \in \mathbb{R}^3$, $i \in [1, N]$ is given by $\mathbf{v}_i = [x_{3i-2}, x_{3i-1}, x_{3i}]^T$. We assume that the meshes have been symmetrically remeshed. This means that for any vertex i , its symmetric partner is

given by $\text{sym}(i)$. Vertices lying on the symmetry line are self-symmetric, i.e. $i = \text{sym}(i)$.

Therefore, each mesh is composed of $N = 2M + S$ vertices, S of which are self-symmetric, leaving $M = (N - S)/2$ pairs of proper symmetric vertices. Without loss of generality, we assume that the ordering of vertices is such that the proper-symmetric vertices on one side of the mesh (e.g. left) come first, followed by the self-symmetric vertices and finally the proper-symmetric vertices on the other (e.g. right) side. Hence:

$$\mathbf{x} = \begin{bmatrix} \mathbf{x}_{\text{left}} \\ \mathbf{x}_{\text{self}} \\ \mathbf{x}_{\text{right}} \end{bmatrix},$$

with

$$\mathbf{x}_{\text{left}} = \mathbf{S}_{\text{left}}\mathbf{x}, \quad \mathbf{x}_{\text{self}} = \mathbf{S}_{\text{self}}\mathbf{x}, \quad \mathbf{x}_{\text{right}} = \mathbf{S}_{\text{right}}\mathbf{x},$$

where

$$\mathbf{S}_{\text{left}} = [\mathbf{I}_{3M} \quad \mathbf{0}_{3M \times 3(M+S)}] \in \{0, 1\}^{3M \times 3N},$$

$$\mathbf{S}_{\text{right}} = [\mathbf{0}_{3M \times 3(M+S)} \quad \mathbf{I}_{3M}] \in \{0, 1\}^{3M \times 3N},$$

$$\mathbf{S}_{\text{self}} = [\mathbf{0}_{3S \times 3M} \quad \mathbf{I}_{3S} \quad \mathbf{0}_{3S \times 3M}] \in \{0, 1\}^{3S \times 3N}$$

are selection matrices that select the proper symmetric vertices from the left and right halves of the mesh and the self symmetric vertices respectively. The vectors \mathbf{x}_{left} and $\mathbf{x}_{\text{right}}$ are assumed to appear in symmetry pair order and so the symmetry operator has a very simple form:

$$\text{sym}(i) = \begin{cases} i + M + S & \text{if } 1 \leq i \leq M \\ i & \text{if } M + 1 \leq i \leq M + S \\ i - M - S & \text{if } M + S + 1 \leq i \leq N \end{cases}$$

A symmetrised mesh is one with exact extrinsic symmetry. For clarity, when we refer to a mesh that has been symmetrised, we use a tilde, $\tilde{\mathbf{x}}$. We assume (without loss of generality) that such symmetrised meshes have their symmetry planes aligned with the $x = 0$ plane. A single vertex is reflected about the $x = 0$ plane by premultiplying with the reflection matrix:

$$\mathbf{F} = \begin{bmatrix} -1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}.$$

so that $\mathbf{F}\mathbf{v}$ is the reflection of \mathbf{v} . This can be extended to the whole template by

$$\mathbf{G}(N) = \mathbf{I}_N \otimes \mathbf{F}$$

so that $\mathbf{G}(N)\mathbf{x}$ is the reflection of \mathbf{x} (we use \otimes to denote the Kronecker product).

To avoid redundancy in symmetrised meshes, we store only the vertices on one side of the mesh since those on the other can be reconstructed by reflection. Moreover, we need only store the y and z coordinates for the self-symmetric vertices since, by definition, their x coordinate is zero. Hence, if $\tilde{\mathbf{x}}$ is a symmetrised mesh, we select the non-redundant entries by:

$$\mathbf{x}_{\text{sym}} = \begin{bmatrix} \mathbf{S}_{\text{left}} \\ \mathbf{S}_{\text{self},yz} \end{bmatrix} \tilde{\mathbf{x}} \in \mathbb{R}^{3M+2S}$$

where

$$\mathbf{S}_{\text{self},yz} = \begin{bmatrix} \mathbf{0}_{2S \times 3M} & \mathbf{I}_S \otimes \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} & \mathbf{0}_{2S \times 3M} \end{bmatrix}$$

is a selection matrix that selects only the y and z components of the self-symmetric vertices.

We can reconstruct a complete mesh from its reduced representation \mathbf{x}_{sym} using:

$$\tilde{\mathbf{x}} = \mathbf{T}\mathbf{x}_{\text{sym}} \quad (1)$$

where

$$\mathbf{T} = \begin{bmatrix} \mathbf{I}_{3M+3S} \\ \mathbf{G}(M)\mathbf{S}_{-\text{self}} \end{bmatrix} \begin{bmatrix} \mathbf{I}_{3M} & \mathbf{0}_{3M \times 2S} \\ \mathbf{0}_{3S \times 3M} & \mathbf{I}_S \otimes \begin{bmatrix} 0 & 0 \\ 1 & 0 \\ 0 & 1 \end{bmatrix} \end{bmatrix}$$

and

$$\mathbf{S}_{-\text{self}} = [\mathbf{I}_{3M} \quad \mathbf{0}_{3M \times 3S}] \in \{0, 1\}^{3M \times 3(M+S)}$$

is a selection matrix that removes the self-symmetric vertices.

5. Proposed pipeline

The following five subsections describe our pipeline in each of the five stages outlined in Sec. 3.

5.1. Symmetry detection and remeshing

The symmetry detection and remeshing can be done simultaneously using template morphing methods. The template we use consists of three parts: left head, symmetry contour and right head. We apply the template morphing method in Hang et al. [9]. This hierarchical parts-based template morphing framework is able to avoid over-fitting and under-fitting and these two properties ensure that the symmetry detection and remeshing is reliable.

5.2. Symmetry plane alignment

For a given mesh, we wish to find a best fit symmetry plane that maximises bilateral extrinsic symmetry. Denote a plane with equation:

$$ax + by + cz = d \quad (2)$$

by $\mathbf{p} = [a, b, c, d]$. We denote by $\text{Ref}_{\mathbf{p}} : \mathbb{R}^3 \mapsto \mathbb{R}^3$ the reflection of a point about the plane given by \mathbf{p} . We seek the plane which satisfies the following optimisation problem:

$$\arg \min_{\mathbf{p}} \sum_i \|\text{Ref}_{\mathbf{p}}(\mathbf{v}_i) - \mathbf{v}_{\text{sym}(i)}\|^2. \quad (3)$$

This problem can be solved in closed form in a straightforward manner. First we compute a modified point cloud in which we take the average of each vertex and its symmetric partner. This gives an almost planar point cloud (it would be exactly planar for a mesh with exact extrinsic symmetry). Second, we apply PCA to this modified point cloud and select the eigenvector with the smallest eigenvalue to give the plane normal. The centre of mass gives a point on the plane. Finally, having found the symmetry plane \mathbf{p} , compute a rigid transformation $[\mathbf{R}, \mathbf{t}]$ such that the symmetry plane coincides with the $x = 0$ plane.

5.3. Symmetrisation

We pose symmetrisation as finding the symmetric mesh that minimises distortion relative to a given mesh, \mathbf{x} , where distortion is measured via linear operators $\mathbf{M}_a, \mathbf{M}_b \in \mathbb{R}^{3N \times 3N}$. We write this as a linear least squares optimisation problem:

$$\arg \min_{\mathbf{x}_{\text{sym}} \in \mathbb{R}^{3M+2S}} \left\| \mathbf{Q} \mathbf{x}_{\text{sym}} - \begin{bmatrix} \mathbf{M}_b \mathbf{x} \\ 0 \\ 0 \end{bmatrix} \right\|^2, \quad (4)$$

where

$$\mathbf{Q} = \begin{bmatrix} \mathbf{M}_a \mathbf{T} \\ \mathbf{1}_{1 \times M} \otimes \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \\ \mathbf{1}_{1 \times S} \otimes \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \end{bmatrix}. \quad (5)$$

The bottom two rows of \mathbf{Q} resolve translational ambiguities by setting the zeroeth moment to zero.

This provides quite a general formulation of the symmetrisation problem. If $\mathbf{M}_a = \mathbf{M}_b = \mathbf{I}_{3N}$ then the above simply performs linear averaging of the left and right halves of the mesh, i.e. it minimises Euclidean distance between the symmetric and original mesh. If instead we construct \mathbf{M}_a and \mathbf{M}_b from the cotangent Laplacian matrix $\mathbf{L} \in \mathbb{R}^{N \times N}$ as

$$\mathbf{M}_{i,j} = \mathbf{L}_{\lceil i/3 \rceil, (j-1 \bmod N)+1}$$

such that $\text{vec}(\mathbf{L}\mathbf{X}) = \mathbf{M} \cdot \text{vec}(\mathbf{X})$, then it minimises local bending distortion relative to the original mesh.

5.4. Symmetry-aware GPA

In Sec. 5.2, we computed a best fit symmetry plane for all meshes. This determines three dimensions of the six degrees of freedom needed to align the meshes to a common

mean (three out of seven if we also allow scaling). Supposing that all faces have been aligned such that their symmetry plane is equal to the $x = 0$ plane then the remaining degrees of freedom are: scaling, rotation about the x axis and translation in y and z .

Denote by $\mathbf{X} \in \mathbb{R}^{3 \times N}$ a training mesh aligned to the $x = 0$ plane and by $\bar{\mathbf{X}} \in \mathbb{R}^{3 \times N}$ the current estimate of the mean (initialised using one of the samples). The Procrustes alignment of a sample to the mean is given by a translation $[0, t_y, t_z]$, a rotation:

$$\mathbf{R}_x(\theta) = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos \theta & -\sin \theta \\ 0 & \sin \theta & \cos \theta \end{bmatrix}, \quad (6)$$

and an optional scale s . To factor out translation, all samples (and the mean) have their y, z centre of mass subtracted:

$$t_y = -\frac{1}{N} \sum_{i=1}^N \mathbf{X}_{2,i}, \quad (7)$$

$$t_z = -\frac{1}{N} \sum_{i=1}^N \mathbf{X}_{3,i}, \quad (8)$$

Ignoring scale for now and assuming that \mathbf{X} and $\bar{\mathbf{X}}$ have had their y, z centre of mass subtracted, we solve the following optimisation problem in terms of the angle of rotation:

$$\begin{aligned} & \arg \min_{\theta} \|\mathbf{R}_x(\theta) \mathbf{X} - \bar{\mathbf{X}}\|_{\text{Fro}}^2 \\ & = \arg \min_{\theta} \sum_{i=1}^N (\cos \theta \mathbf{X}_{2,i} - \sin \theta \mathbf{X}_{3,i} - \bar{\mathbf{X}}_{2,i})^2 \\ & \quad + (\sin \theta \mathbf{X}_{2,i} + \cos \theta \mathbf{X}_{3,i} - \bar{\mathbf{X}}_{3,i})^2 \end{aligned}$$

Differentiating with respect to θ , setting to zero and solving for θ yields:

$$\begin{aligned} \mathbf{F}_1 &= \sum_{i=1}^N (\mathbf{X}_{2,i} \bar{\mathbf{X}}_{3,i} - \mathbf{X}_{3,i} \bar{\mathbf{X}}_{2,i}), \\ \mathbf{F}_2 &= \sum_{i=1}^N (\mathbf{X}_{2,i} \bar{\mathbf{X}}_{2,i} + \mathbf{X}_{3,i} \bar{\mathbf{X}}_{3,i}), \\ \theta &= \text{atan2}(\mathbf{F}_1, \mathbf{F}_2). \end{aligned}$$

Having aligned all of the meshes to the mean using the above steps, we compute a new mean and iterate. We refer to this process as Symmetry-aware GPA (SGPA).

5.5. Building a symmetry-factored model

We now build a statistical model composed of two parts: a symmetric part (which only need model points on one side of the symmetry plane and the self-symmetric points on the

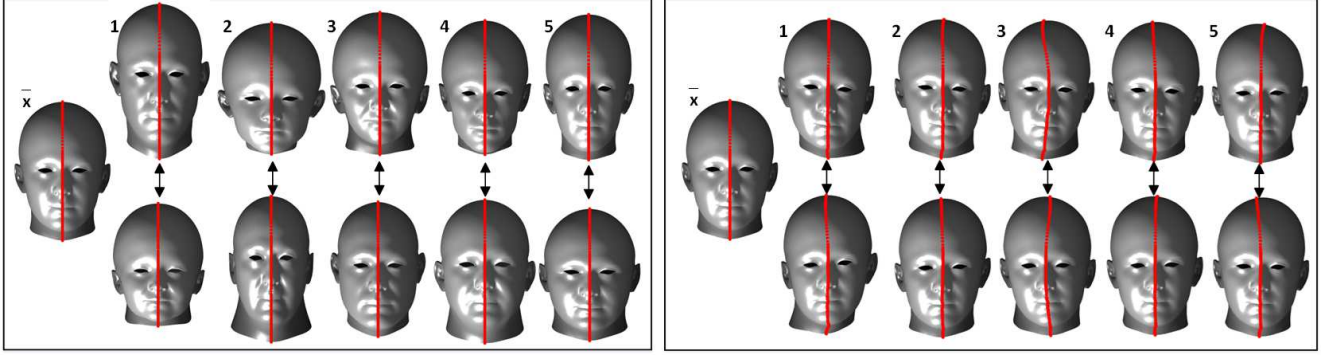


Figure 1. Visualisation of 5 principal components (mean ± 5 SDs): Left box - Symmetry variation ; Right box - Asymmetry variation

symmetry plane, since the other half is by definition given by a reflection), and an asymmetric part (modelled as displacements to the output of the symmetric model).

We now construct a symmetric data matrix $\mathbf{D}_{\text{sym}} \in \mathbb{R}^{(3M+2S) \times K}$ from the K symmetrised and SGPA-aligned training meshes. We apply PCA to find eigenvectors/eigenvalues of the symmetrised meshes. The symmetric statistical model is given by:

$$\mathbf{x}_{\text{sym}} = \bar{\mathbf{x}}_{\text{sym}} + \mathbf{P}_{\text{sym}} \mathbf{b}_{\text{sym}}, \quad (9)$$

where $\mathbf{P}_{\text{sym}} \in \mathbb{R}^{(3M+2S) \times K}$, $\bar{\mathbf{x}}_{\text{sym}} \in \mathbb{R}^{3M+2S}$. A complete symmetric mesh can be built from \mathbf{x}_{sym} using (1).

The data matrix for the asymmetric model, $\mathbf{D}_{\text{asym}} \in \mathbb{R}^{3N \times K}$, is constructed by subtracting the symmetrised meshes from the original meshes, such that the i th column is given by:

$$\mathbf{D}_{\text{asym},i} = \mathbf{x}_i - \mathbf{T}\mathbf{D}_{\text{sym},i}. \quad (10)$$

We again apply PCA to this matrix but without computing and subtracting a mean since the data matrix directly describes displacements. The final combined model is hence given by:

$$\mathbf{x} = \mathbf{T}(\bar{\mathbf{x}}_{\text{sym}} + \mathbf{P}_{\text{sym}} \mathbf{b}_{\text{sym}}) + \mathbf{P}_{\text{asym}} \mathbf{b}_{\text{asym}} \quad (11)$$

Note that this could be rewritten in the form:

$$\mathbf{x} = \bar{\mathbf{x}} + \mathbf{P} \begin{bmatrix} \mathbf{b}_{\text{sym}} \\ \mathbf{b}_{\text{asym}} \end{bmatrix} \quad (12)$$

where

$$\bar{\mathbf{x}} = \mathbf{T}\bar{\mathbf{x}}_{\text{sym}} \quad (13)$$

$$\mathbf{P} = \begin{bmatrix} \mathbf{P}_{\text{sym}} & \mathbf{P}_{\text{asym}} \\ \mathbf{G}(M)\mathbf{S}_{\text{-self}}\mathbf{P}_{\text{sym}} & \mathbf{0} \end{bmatrix}. \quad (14)$$

Hence, it is still just a standard linear model but for which the parameters can be partitioned into those that only vary the shape symmetrically and those that vary it asymmetrically.

6. Experiments

We use 1212 individuals (606 males and 606 females) [9, 11] to derive our symmetry-factored 3D craniofacial model, by applying the proposed construction pipeline. We compare the proposed model with linear PCA models which are symmetry-aware GPA + PCA and standard GPA + PCA.

6.1. 3D morphable model of asymmetry

As can be seen from Fig.1, the left box shows the morphable model of symmetry variation: mean ± 5 standard deviations of symmetry variation; the right box shows the morphable model of asymmetry variation. In order to observe pure asymmetry variation, we use mean of symmetry model ± 5 standard deviations of asymmetry variation. Red points are the symmetry contour. Note that the major mode of asymmetry is the angle between head and neck.

6.2. Model fitting

Fig.2 illustrates the end results of symmetry plane alignment, symmetrisation and model fitting. After symmetrisation, the distance error color map becomes symmetrical, see Fig.2 (2). The symmetry plane of Fig.2 (2) is the $x = 0$ plane. The aim of model fitting is to find the parameters \mathbf{P} in Eqn.12 and validate the effectiveness of morphable models. Figure 3 shows the cumulative error distributions of the per-vertex fitting error. For very small fitting errors (less than 0.487 mm), the standard PCA model is better than the proposed model; but for larger fitting errors, which is 91.3% of the dataset, the proposed model is better than the linear PCA models. The proposed symmetry-aware GPA improves the performance of the standard PCA model.

6.3. Model evaluation

We compute model compactness and generalisation error in order to evaluate the proposed model quantitatively. To ensure a fair comparison, when we compute these two criteria, the number of principal components (PCs) from the

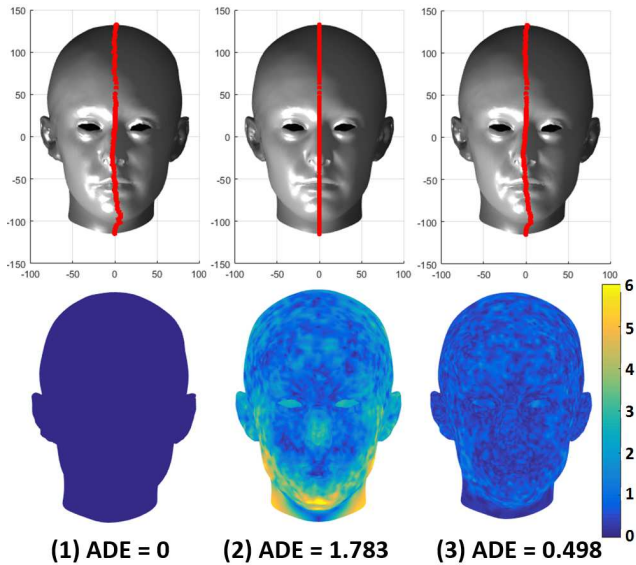


Figure 2. (1) Symmetry plane alignment, red line shows deformed symmetry contour from initial template morphing (2) symmetrisation, note the straightened symmetry contour (3) model fitting; Average distance error (ADE, mm) against the shape sample (1).

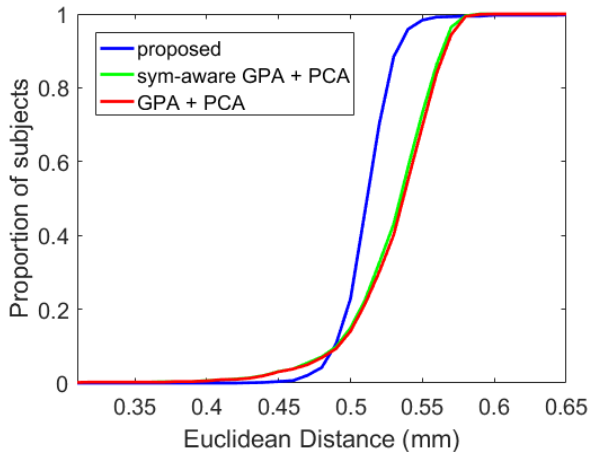


Figure 3. Cumulative error distributions of the per-vertex fitting error (mm): proposed V.S. standard PCA model (higher is better)

standard PCA model is the sum of PCs from symmetry-factored model, i.e. if we use N_{pc} PCs from the standard PCA model, the number of PCs from the symmetry model is $\frac{N_{pc}}{2}$ and also $\frac{N_{pc}}{2}$ from the asymmetry model. As can be seen from Fig.4, the asymmetry model requires fewer components to express its variation than the symmetry model, which in turn requires fewer components than the PCA model. When we combine $\frac{N_{pc}}{2}$ symmetry and $\frac{N_{pc}}{2}$ asymmetry components with N_{pc} PCA components in the compactness plot in Fig. 5, we find that the proposed model is slightly better than the standard PCA model and more compact than the symmetry-aware GPA + PCA model.

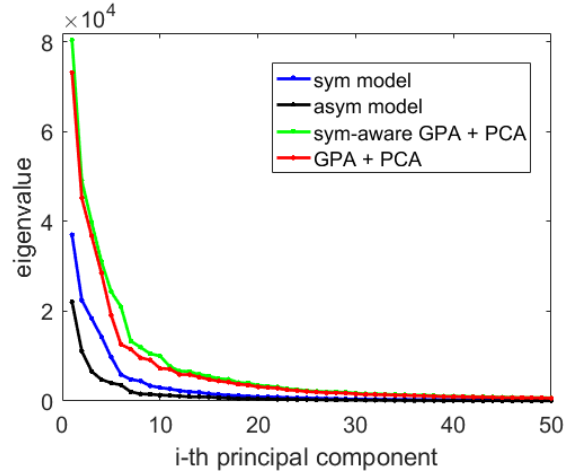


Figure 4. Plot of eigenvalues. Fewer PCs describe more variation

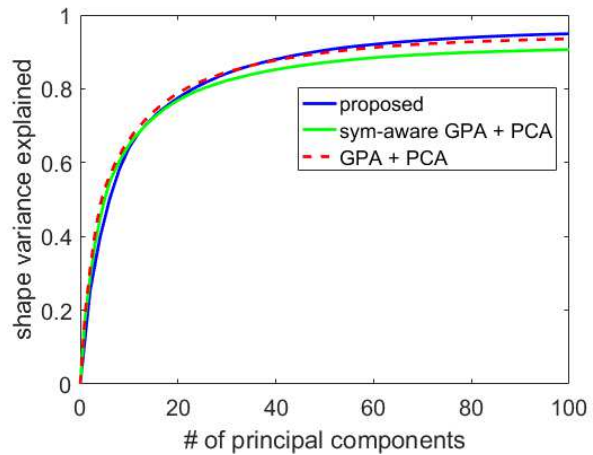


Figure 5. Compactness, higher is better

The generalization ability of a model measures its capability to represent unseen instances of the object class [35]. The generalization ability of each model is measured using leave-one-out reconstruction. A model is built using all but one member of the training set and then fitted to the excluded example. The generalisation error shown in Fig.6 implies that the proposed model has significantly better generalisation ability than the linear PCA models. The symmetry-aware GPA + PCA model has slightly lower reconstruction error than the standard PCA model, which implies that the proposed symmetry-aware GPA can improve model generalization ability.

6.4. Half head completion

When shade-from-shading is applied to a 2D profile image, the standard way to make the depth information complete is to calculate the reflection of the half head. There is no asymmetry variation in this completion. A PCA model can retain asymmetry variation in completion, but it tends




Figure 6. Generalisation error of the proposed model and PCA model with the number of principal components retained (lower is better).

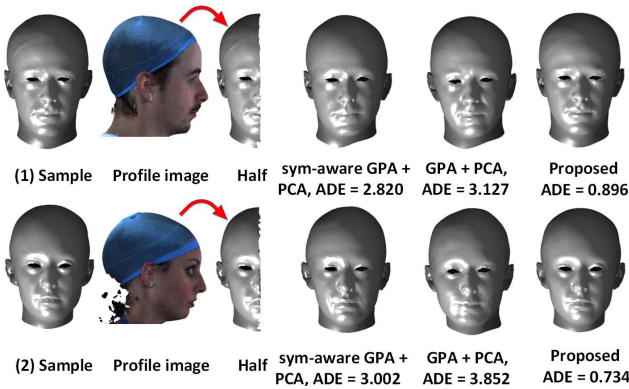


Figure 7. Simulation of half head completion from profile image.

to lose more symmetry information. The proposed model is able to overcome the loss of both symmetry and asymmetry information. From Eqn.12, the reconstruction from the proposed model can be decomposed into two steps:

- (i) Find the closest symmetry head, computing the symmetry parameters \mathbf{P}_{sym} by least squares;
- (ii) Use \mathbf{P}_{sym} to find the closest asymmetry variation, computing the asymmetry parameters \mathbf{P}_{asym} by least squares.

Step (i) retains symmetry information and step (ii) retains asymmetry information. This can be validated by the results in Fig.7 and Fig.8 in half-head completion experiments. We use 606 training samples to build the 3DMMs of craniofacial asymmetry and use the rest of dataset (606 subjects) for the half-head completion task. 200 PCs from PCA model are used to do the half head completion. For a fair comparison, 100 PCs from symmetry model and 100 PCs from asymmetry model are used. As shown in Fig.8, the pro-

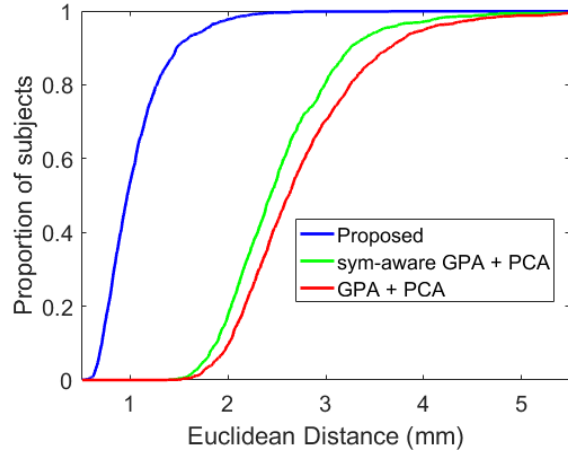


Figure 8. Cumulative error distributions of the per-vertex fitting error (mm) from half head: proposed model V.S. standard PCA model (higher is better)

Table 1. Age classification results

	Precision	Recall	F-score
PCA	0.733	0.735	0.731
Sym	0.737	0.741	0.736
Asym	0.709	0.712	0.710
Sym+Asym	0.739	0.741	0.741

posed model has much greater proportion of subjects with lower reconstruction error from half head.

6.5. Age classification: asymmetry V.S. symmetry

We divide the dataset into four age groups: age 0-11, 12-21, 22-60 and > 60 . \mathbf{P}_{sym} and \mathbf{P}_{asym} are used for the age classification task and compared with the standard PCA model. The comparison uses

- the first 10 parameters from the PCA model
- the first 10 from \mathbf{P}_{sym}
- the first 10 from \mathbf{P}_{asym}
- the first 5 from \mathbf{P}_{sym} and 5 from \mathbf{P}_{asym}

A linear SVM is trained with 10-fold cross-validation. As can be seen from the Table.1, symmetry parameters have slightly better accuracy than that of PCA model, while asymmetry parameters alone are worse in age classification than that of PCA model. This implies that symmetry information is more discriminative in age classification than asymmetry information. Since the asymmetry variation is relevant to the age (faces in age group 0-11 are more symmetrical than that in > 60 age group), the asymmetry parameters are still valid in age classification. The combination of 5 \mathbf{P}_{sym} and 5 \mathbf{P}_{asym} are better in age classification than PCA model and \mathbf{P}_{sym} alone, which implies that the

asymmetry information can improve the performance in age classification when combined with symmetry parameters.

7. Conclusion

We presented a craniofacial symmetry-factored statistical modelling pipeline and applied it to a craniofacial dataset of 1212 subjects. Contributions included Laplace-Beltrami regularized symmetrisation, symmetry-aware GPA and symmetry-factored statistical modelling method. We used it to build the first morphable model that makes craniofacial asymmetry explicit. A comprehensive evaluation shows that the proposed model has significantly better performance than the linear PCA models, especially in terms of generalisation error and in the completion of head data. Symmetry-aware GPA can improve the performance of the standard PCA model. Additionally, we found that symmetry information is more discriminative than asymmetry information in age classification. Future work will focus on region-based asymmetry modelling and give a comprehensive study of facial symmetry/asymmetry in 3D face recognition.

Acknowledgements

We thank Google Faculty Awards, the Royal Academy of Engineering and the Leverhulme Trust for their support. Headspace data collection was supported by QIDIS from the National Commissioning Group. We also thank Rachel Armstrong, Headspace project coordinator.

References

- [1] B. Amberg, S. Romdhani, and T. Vetter. Optimal step non-rigid icp algorithms for surface registration. In *Proceedings of CVPR*, pages 1–8, 2007. 2
- [2] V. Blanz and T. Vetter. A morphable model for the synthesis of 3d faces. In *Proceedings of Computer graphics and interactive techniques*, pages 187–194, 1999. 2
- [3] V. Blanz and T. Vetter. Face recognition based on fitting a 3d morphable model. *IEEE Transactions on pattern analysis and machine intelligence*, 25(9):1063–1074, 2003. 2
- [4] J. Booth, A. Roussos, S. Zafeiriou, A. Ponniah, and D. Dunaway. A 3d morphable model learnt from 10,000 faces. In *Proceedings of CVPR*, pages 5543–5552, 2016. 2
- [5] J. Chen, C. Yang, Y. Deng, G. Zhang, and G. Su. Exploring facial asymmetry using optical flow. *IEEE Signal Processing Letters*, 21(7):792–795, 2014. 2
- [6] P. Claes, M. Walters, D. Vandermeulen, and J. G. Clement. Spatially-dense 3d facial asymmetry assessment in both typical and disordered growth. *Journal of anatomy*, 219(4):444–455, 2011. 2
- [7] B. Combès, R. Hennessy, J. Waddington, N. Roberts, and S. Prima. Automatic symmetry plane estimation of bilateral objects in point clouds. In *CVPR 2008*, pages 1–8, 2008. 2
- [8] D. Cosker, E. Krumhuber, and A. Hilton. A faces valid 3d dynamic action unit database with applications to 3d dynamic morphable facial modeling. In *Proceedings of ICCV*, pages 2296–2303, 2011. 2
- [9] H. Dai, N. Pears, W. Smith, and C. Duncan. A 3d morphable model of craniofacial shape and texture variation. In *Proceedings of ICCV*, 2017. 1, 2, 3, 5
- [10] M. De Smet and L. Van Gool. Optimal regions for linear model-based 3d face reconstruction. In *Proceedings of Asian Conference on Computer Vision*, pages 276–289, 2010. 2
- [11] C. Duncan, H. Dai, N. Pears, and W. Smith. The liverpool-york head model a novel shape based outcomes analysis tool for craniofacial surgeons. In *17th Biennial Congress of the International Society of Craniofacial Surgery (ISCFS)*, 2017. 5
- [12] A. Golovinskiy, J. Podolak, and T. Funkhouser. Symmetry-aware mesh processing. *Mathematics of Surfaces XIII*, pages 170–188, 2009. 1
- [13] R. J. Hennessy, S. McLearie, A. Kinsella, and J. L. Waddington. Facial shape and asymmetry by three-dimensional laser surface scanning covary with cognition in a sexually dimorphic manner. *The Journal of neuropsychiatry and clinical neurosciences*, 18(1):73–80, 2006. 2
- [14] B. C. Jones, A. C. Little, I. S. Penton-Voak, B. Tiddeman, D. M. Burt, and D. Perrett. Facial symmetry and judgements of apparent health: support for a good genes explanation of the attractiveness–symmetry relationship. *Evolution and human behavior*, 22(6):417–429, 2001. 1
- [15] J. T. Kent and K. V. Mardia. Shape, procrustes tangent projections and bilateral symmetry. *Biometrika*, pages 469–485, 2001. 2
- [16] C. P. Klingenberg, M. Barluenga, and A. Meyer. Shape analysis of symmetric structures: quantifying variation among individuals and asymmetry. *Evolution*, 56(10):1909–1920, 2002. 2
- [17] R. R. Kolamunnage-Dona and J. T. Kent. Decomposing departures from bilateral symmetry. In S. Barber, P. Baxter, K. V. Mardia, and R. E. Walls, editors, *Quantitative Biology, Shape Analysis, and Wavelets*, pages 75–78. Leeds University Press, 2005. 2
- [18] Q. Liao, X. Jin, and W. Zeng. Enhancing the symmetry and proportion of 3d face geometry. *IEEE transactions on visualization and computer graphics*, 18(10):1704–1716, 2012. 2
- [19] Y. Liu and J. Palmer. A quantified study of facial asymmetry in 3d faces. 2003. 2
- [20] K. V. Mardia, F. L. Bookstein, and I. J. Moreton. Statistical assessment of bilateral symmetry of shapes. *Biometrika*, pages 285–300, 2000. 2
- [21] D. Milner, S. Raz, H. Hel-Or, D. Keren, and E. Nevo. A new measure of symmetry and its application to classification of bifurcating structures. *Pattern Recognition*, 40(8):2237–2250, 2007. 2
- [22] N. J. Mitra, L. J. Guibas, and M. Pauly. Partial and approximate symmetry detection for 3d geometry. In *ACM Transactions on Graphics (TOG)*, volume 25, pages 560–568. ACM, 2006. 1

- [23] N. J. Mitra, L. J. Guibas, and M. Pauly. Symmetrization. *ACM Transactions on Graphics (TOG)*, 26(3):63, 2007. 2
- [24] N. J. Mitra, M. Pauly, M. Wand, and D. Ceylan. Symmetry in 3d geometry: Extraction and applications. In *Computer Graphics Forum*, volume 32, pages 1–23, 2013. 2
- [25] M. Ovsjanikov, J. Sun, and L. Guibas. Global intrinsic symmetries of shapes. In *Computer graphics forum*, volume 27, pages 1341–1348. Wiley Online Library, 2008. 1
- [26] G. Pan, Y. Wang, Y. Qi, and Z. Wu. Finding symmetry plane of 3d face shape. In *Pattern Recognition, 2006. ICPR 2006. 18th International Conference on*, volume 3, pages 1143–1146, 2006. 2
- [27] G. Passalis, P. Perakis, T. Theoharis, and I. A. Kakadiaris. Using facial symmetry to handle pose variations in real-world 3d face recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 33(10):1938–1951, 2011. 2
- [28] V. Patraucean, R. Grompone von Gioi, and M. Ovsjanikov. Detection of mirror-symmetric image patches. In *Proc. CVPR Workshops*, pages 211–216, 2013. 1
- [29] P. Paysan, R. Knothe, B. Amberg, S. Romdhani, and T. Vetter. A 3d face model for pose and illumination invariant face recognition. In *Proceedings of Advanced video and signal based surveillance.*, pages 296–301, 2009. 2
- [30] J. Podolak, P. Shilane, A. Golovinskiy, S. Rusinkiewicz, and T. Funkhouser. A planar-reflective symmetry transform for 3d shapes. *ACM Transactions on Graphics (TOG)*, 25(3):549–559, 2006. 1
- [31] J. Roth, Y. Tong, and X. Liu. Adaptive 3d face reconstruction from unconstrained photo collections. In *Proceedings of CVPR*, pages 4197–4206, 2016. 2
- [32] A. Salazar, S. Wuhler, C. Shu, and F. Prieto. Fully automatic expression-invariant face correspondence. *Machine Vision and Applications*, 25(4):859–879, 2014. 2
- [33] Y. Savriama and C. P. Klingenberg. Beyond bilateral symmetry: geometric morphometric methods for any type of symmetry. *BMC evolutionary biology*, 11(1):280, 2011. 2
- [34] A. Shehu, A. Brunton, S. Wuhler, and M. Wand. *Characterization of Partial Intrinsic Symmetries*, pages 267–282. 2014. 1
- [35] M. A. Styner, K. T. Rajamani, L.-P. Nolte, G. Zsemlye, G. Székely, C. J. Taylor, and R. H. Davies. Evaluation of 3d correspondence methods for model building. In *Biennial International Conference on Information Processing in Medical Imaging*, pages 63–75, 2003. 6
- [36] L. Zhang, A. Razdan, G. Farin, J. Femiani, M. Bae, and C. Lockwood. 3d face authentication and recognition based on bilateral symmetry analysis. *The Visual Computer*, 22(1):43–55, 2006. 2
- [37] A. I. Zhurov, S. Richmond, C. H. Kau, and A. Toma. Averaging facial images. *Three-dimensional imaging for orthodontics and maxillofacial surgery*. London: Wiley-Blackwell, pages 126–44, 2010. 2